HIERARCHICAL RECONSTRUCTION USING GEOMETRY AND SINOGRAM RESTORATION

Jerry L. Prince and Alan S. Willsky

Affiliations:

Jerry L. Prince:
Image Analysis and Communications Laboratory
Department of Electrical and Computer Engineering
The Johns Hopkins University
Baltimore, MD 21218

Alan S. Willsky:
Laboratory for Information and Decision Systems
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139

Correspond to:

Jerry L. Prince
105 Barton Hall
The Johns Hopkins University
Baltimore, MD 21218
Tel: (410) 516-5192
Fax: (410) 516-5566
E-mail: prince@mashie.ece.jhu.edu

IP Editors' Information Classification Scheme (EDICS):

2.3

Report Documentation Page		Form Approved OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **1991**	2. REPORT TYPE	3. DATES COVERED **00-00-1991 to 00-00-1991**
4. TITLE AND SUBTITLE **Hierarchical Reconstruction Using Geometry and Sinogram Restoration**		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute of Technology,Laboratory for Information and Decision Systems,77 Massachusetts Avenue,Cambridge,MA,02139-4307**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES **52**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

HIERARCHICAL RECONSTRUCTION USING GEOMETRY AND SINOGRAM RESTORATION

Jerry L. Prince and Alan S. Willsky

ABSTRACT

We describe and demonstrate a hierarchical reconstruction algorithm for use in noisy and limited-angle or sparse-angle tomography. The algorithm estimates an object's mass, center of mass, and convex hull from the available projections, and uses this information, along with fundamental mathematical constraints, to estimate a full set of smoothed projections. The mass and center of mass estimates are made using a least squares estimator derived from the principles of consistency of the Radon transform. The convex hull estimate is produced by first estimating the positions of support lines of the object from each available projection and then estimating the overall convex hull using maximum likelihood techniques or maximum *a posteriori* techniques. Estimating the position of two support lines from a single projection is accomplished using a generalized likelihood ratio technique for estimating jumps in linear systems. We show results for simulated objects in a variety of measurement situations and for several model parameters and discuss several possible extensions to the work.

J.L. Prince is with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218.

A.S. Willsky is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

I. INTRODUCTION

Computed tomography (CT), the practice of reconstructing images of cross-sections from measurements of their projections, has become an important tool in many areas of application, including non-destructive evaluation, sonar imaging, synthetic aperture imaging, and medical imaging. It is well known that in the full-data problem, where one is given enough high-quality projections over a 180 degree angular range, images of outstanding quality may be obtained. In many important practical cases, however, there is not enough data to obtain high-quality images using the usual techniques. In particular, the *limited-angle problem* occurs when projections are available over an angular range less than 180 degrees, and the *sparse-angle problem* occurs when only a small number of angles evenly spaced over 180 degrees are available.

Limited angle and sparse and problems occur often in practice. For example, in medical imaging the data acquisition for stop-action cardiac CT imaging must occur so rapidly that the carriage containing the X-ray emitters and detectors can only travel part of the way through the full angular range before significant heart motion occurs [1]. Another example, from non-destructive testing, is the tomographic suitcase scanner for airport security [2], in which a fan-beam X-ray emitter is placed above the suitcase with a set of detectors positioned below. Line integrals that are nearly horizontal cannot be measured in this case. These are both limited-angle problems. Sparse-angle problems often arise when projection methods are used in fluid dynamics and combustion analysis. For example, Zoltani et. al. [3] point out that in the ballistic environment, where one one wishes to analyze two-phase flow under actual firing conditions, the time and energy constraints are so severe that acquiring a sparse data set is unavoidable.

Noise is another problem that causes degradation of reconstructed tomographic images. Noise arises in different ways depending on the application. For example, in applications where X-rays are used, if the X-rays have low energy or the sample has high attenuation, the noise is dominated by the photon statistics of the X-rays [4]. In synthetic aperture radar problems, receiver noise, clutter, and jamming all contribute to the overall noise that obscures the desired signal [5].

Many researchers have proposed solutions to the limited-angle and sparse-angle problems, although few have also dealt explicitly with noise. Solutions tend to fall into two categories: transform techniques that incorporate little *a priori* information and finite series expansion methods that may incorporate *a priori* information as constraints or probabilities. The transform techniques are usually single-pass direct reconstructions while the finite series expansion

methods are usually iterative. Among the transform techniques are the pseudoinverse of the 2-D Radon transform [6, 7], angle-dependent rho-filters [8], analytic continuation in the Fourier plane [9], and the method of squashing [10]. The finite series expansion methods include linear minimum variance methods [11, 12], projection onto convex sets (POCS) [13, 14], the Gerchberg-Papoulis algorithm [15, 16], iterative pseudoinverse methods [14, 17], and Bayesian methods [18, 19].

In addition to these these two general approaches there are other approaches that depend upon severely restricting the class of objects to be reconstructed. For example, Rossi and Willsky [20] use hierarchical maximum likelihood methods to estimate the position, radius, and eccentricity of objects with a known unit profile such as the unit disk. Soumekh [21], Chang and Shelton [22], and Fishburn et al. [23] have investigated reconstruction of binary objects from a small number of projections.

The hierarchical algorithm described in this paper was designed to address several problems common to many of the existing algorithms mentioned above. For example, one of the major problems in methods that iterate between the object and projection spaces is reprojection error [24]. This problem is particularly bad when noise is present in the measurements [12]. One elegant way to avoid this problem, proposed by Kim, Kwak, and Park [25, 26], is to iterate entirely in projection-space, incorporating certain object constraints mathematically, without reprojection. Our method is also a projection-space method, although it is not a POCS method as in [25] and [26].

Another problem with many, but not all, of the existing limited-angle algorithms is that they do not account for noise in an optimal sense. Although it is possible to modify POCS to account for noise processes [27], it does not make optimal (in a Bayesian sense) use of known noise statistics together with *a priori* knowledge. Minimum variance and Bayesian methods do take noise processes into account, but rarely have the capability to account for additional geometric knowledge as is done in POCS and our method. Parts of our method are Bayesian also, using the maximum *a posteriori* criterion, which specifies optimum solutions accounting for both the noise process and *a priori* knowledge.

Finally, a vital part of many POCS algorithms, the Gerchberg-Papoulis algorithm, and some iterative pseudoinverse approaches is the necessity of having known convex constraints, and in particular, having knowledge of the object support or convex support. Medoff [14] notes that one way to acquire this information is to have "a radiologist determine the outer boundary of the object". Aside from the fact that this is a potentially time-consuming process for a busy radiologist, the image that the radiologist uses to generate this boundary must

be created before any correction has taken place, and therefore can be expected to be rife with artifacts from both noise and limited-angle measurement geometry. One has therefore introduced a potential source of error in this seemingly innocuous step. Our method, instead, estimates the convex support of the object directly from the projection measurements using prior probabilistic knowledge of the general shape to be expected. We then use this estimate to assist in a projection-space reconstruction algorithm.

The algorithm described in this paper uses a hierarchical approach to process the available measured projections in order to generate an estimate of the full set of projections, an image of which is called a *sinogram*. The object estimate is then produced using convolution backprojection applied to this estimated sinogram. The estimated sinogram satisfies the two lowest-order constraints, which we call the mass and center of mass constraints, specified by the Ludwig-Helgason consistency conditions of the Radon transform [28, 29]. It is also consistent with the derived convex support estimate, in that its values are nearly zero for lines that miss the estimated convex support. Finally, it is smoothed by incorporating prior probabilistic knowledge using a Markov random field, optimally removing the contributions of noise.

In [30] we presented an algorithm for sinogram restoration which uses assumed knowledge of the mass, center of mass, and convex support of the object to be reconstructed, together with a prior Markov random field probabilistic description of sinograms, in order to estimate smoothed sinograms from noisy and possibly limited-angle or sparse-angle data. This procedure forms the core of the algorithm described in this paper. The nature of the *a priori* information required by this procedure, however, precludes its use as a stand-alone procedure for processing raw projection data. In this paper we develop methods to estimate the mass, center of mass, and convex support of the object directly from the projections, hence eliminating the requirement that these be known *a priori*, and thus providing a complete sinogram restoration algorithm. To estimate the convex support of the object, we draw upon work appearing in [31] and [32], where we present methods to estimate convex sets from noisy support line measurements. Some of these methods can incorporate prior geometric knowledge about the shape of the convex support, a necessity in the limited-angle and sparse-angle cases, where there is otherwise not enough data to provide unique solutions. In this work, we complete the picture by developing methods to estimate the mass and center of mass, and methods to estimate the lateral positions of support lines of an object directly from the projections. Combining these new methods with those in [30], [31], and [32], and characterizing and utilizing the sources of error generated at each step (in

order to determine automatically several needed algorithm parameters) gives the hierarchical algorithm described herein.

The rest of the paper is organized as follows. Section II reviews the properties of sinograms, particularly the consistency and convex support relations. Section III reviews our previous work on sinogram restoration and convex set estimation. Section IV presents new material: methods for support value estimation, mass and center of mass estimation, error estimation, and the full hierarchical algorithm. Simulation results are presented in Section V, including several limited-angle and sparse-angle scenarios. Finally, Section VI summarizes the significant points and discusses possible future research directions.

II. PROPERTIES OF SINOGRAMS

The 2-D Radon transform is given by [33]

$$g(t,\theta) = \int_{\mathbf{x} \in \mathbf{R}^2} f(\mathbf{x})\delta(t - \boldsymbol{\omega}^{\mathrm{T}}\mathbf{x})d\mathbf{x}, \qquad (1)$$

where $\boldsymbol{\omega} = [\cos\theta \quad \sin\theta]^{\mathrm{T}}$, θ is the angle measured counterclockwise from the horizontal-axis, $\delta(\cdot)$ is the Dirac delta function, and $f(\mathbf{x})$ is a function of the two-dimensional vector $\mathbf{x}$. For this paper, we assume $f(\mathbf{x})$ to be a real function defined on the disk of radius T centered at the origin. Figure 1 shows the geometry of the 2-D Radon transform. For a particular θ the function $g(t,\theta)$ is a function of t, and is called a *parallel ray projection* or just a *projection*. A *sinogram* is an image of the 2-D Radon transform, where t and θ form the vertical and horizontal axes, respectively, of a cartesian coordinate system. Because of the periodicity of the 2-D Radon transform and because of the assumed domain of $f(\mathbf{x})$, the sinogram is completely characterized by knowledge of $g(t,\theta)$ over the domain

$$\mathcal{Y}_T = \{(t,\theta) \mid t \in [-T,T], \theta \in [\pi/2, 3\pi/2)\}. \qquad (2)$$

An object and its sinogram are displayed in Figure 2. Note that the columns of the sinogram are projections, with the left-most projection arising from horizontal line integrals. In most real problems, we expect to have a discrete version of the sinogram, sampled for many values of t and θ. We define a *finest-grain sinogram* to be one that is known over the rectilinear lattice of n_d (odd) uniformly spaced points in the t-direction and n_v uniformly spaced points in the θ-direction. Our observations, both limited-angle and sparse-angle, consist of measured sinogram values over a subset $\mathcal{Y}_O$ of this finest-grain lattice.

A. Mass and Center of Mass

Conventional tomographic reconstruction algorithms, such as convolution backprojection (CBP) [4], which attempt to invert the Radon transform, require the availability of a complete set of projection data in order for the inversion to be possible. Consequently, their use in limited-angle or sparse-angle situations requires that some accommodation be made for the missing data. The simplest approach is in essence to set the missing measurement values to zero by applying the inversion operator only over the available measurement set. Such an approach is well-known to produce a severely distorted reconstruction, and, while simple schemes such as linear interpolation [34] typically lead to some improvement, serious degradations are still present. Similarly, the presence of significant measurement errors in the projection data can lead to pronounced distortions or artifacts in the resulting reconstructions [35]. One of the reasons for the presence and level of severity of degradations in each of these cases is that the inversion operation is being applied to a data set that could *not* be the Radon transform of *any* object. Specifically, a function $g(t, \theta)$ that is a valid 2-D Radon transform must satisfy the Ludwig-Helgason consistency conditions [28, 29], which are summarized in the following theorem.

Theorem 1 (Consistency Theorem) *Let $\mathcal{S}$ denote the space of rapidly decreasing C^∞ function on $\mathbb{R}^2$. In order for $g(t, \theta)$ to be the 2-D Radon transform of a function $f \in \mathcal{S}$, it is necessary and sufficient that*

(a) $g \in \mathcal{S}(\mathbb{R}^1 \times S^1)$,

(b) $g(t, \theta + \pi) = g(-t, \theta)$, and

(c) the integral

$$\int_{-\infty}^{\infty} g(t, \theta) t^k \, dt \tag{3}$$

be a homogeneous polynomial of degree k in $\cos \theta$ and $\sin \theta$ for all $k \geq 0$.

Proof *See [28].*   □

In Theorem 1, condition (b) specifies the sinogram periodicity condition; and the two lowest moments of condition (c), i.e., where $k = 0$ and $k = 1$, give rise to the *mass* and *center of mass* sinogram constraints:

$$m(\theta) = \int_{-\infty}^{\infty} g(t, \theta) dt = m \tag{4}$$

and

$$c(\theta) = \frac{1}{m} \int_{-\infty}^{\infty} t g(t, \theta) dt = \mathbf{c}^{\mathrm{T}} \boldsymbol{\omega} = c_1 \cos \theta + c_2 \sin \theta \,, \tag{5}$$

where

$$m = \int_{\mathbf{x} \in \mathbf{R}^2} f(\mathbf{x}) \, d\mathbf{x} \,, \tag{6}$$

and

$$\mathbf{c} = [c_1 \; c_2]^{\mathrm{T}} = \frac{1}{m} \int_{\mathbf{x} \in \mathbf{R}^2} \mathbf{x} f(\mathbf{x}) \, d\mathbf{x} \,. \tag{7}$$

One of the aspects of our algorithm is to use these consistency constraints in order to guarantee that our sinogram smoothing and interpolation operations yield consistent sinograms, thereby significantly reducing degradations and artifacts caused by noisy or missing data. Although there are an infinite number of constraints which any valid 2-D Radon transform must satisfy, in this paper we focus only on these two lowest order moments. The reason for this is that they have simple geometric interpretations and that they may be easily and reliably estimated from the available projections, as we shall see in Section IV.

B. Convex Support

One very useful way in which to interpret the Radon transform consistency constraints is that they provide prior information that in essence reduces the number of degrees of freedom that must be recovered from the measurement data. A second piece of prior information which can be of significant value for the same reason is knowledge about the support, $\mathcal{F}$, of the function $f(\mathbf{x})$ to be reconstructed, i.e. the set of points where $f(\mathbf{x})$ may be non-zero. In particular, our algorithm makes use of information about the *convex support* of $f(\mathbf{x})$, i.e., the convex hull of $\mathcal{F}$, $\mathcal{F}_c = \mathrm{hul}(\mathcal{F})$. For any set S, the *support value* of S at angle θ is the maximum projection of points in S onto the axis at angle θ with unit vector $\boldsymbol{\omega} = [\cos \theta \; \sin \theta]^{\mathrm{T}}$. This is given by

$$h(\theta) = \sup_{\mathbf{x} \in S} \mathbf{x}^{\mathrm{T}} \boldsymbol{\omega} \,. \tag{8}$$

Treated as a function of θ, $h(\theta)$ is known as the *support function* of θ.

As shown in Figure 3 we see that for a fixed θ, the projection $g(t, \theta)$ has support confined to the interval between the two points

$$t_+(\theta) = \sup_{r \in \{t \,|\, g(t,\theta) \neq 0\}} r \,, \tag{9}$$

$$t_-(\theta) = \inf_{r \in \{t \,|\, g(t,\theta) \neq 0\}} r \,, \tag{10}$$

which are related to the support function as follows:

$$h(\theta) = \begin{cases} t_+(\theta), & 0 \le \theta < \pi \\ -t_-(\theta - \pi), & \pi \le \theta < 2\pi \end{cases}. \qquad (11)$$

As shown in Figure 4, the collection of support values, or equivalently the **support** function, segment a sinogram into a region of support $\mathcal{G}$ for $g(t,\theta)$ and its **complement** $\bar{\mathcal{G}}$, where

$$\mathcal{G} = \{(t,\theta) \in \mathcal{Y}_T \mid t_-(\theta) \le t \le t_+(\theta)\}. \qquad (12)$$

Therefore, for a given object support set $\mathcal{F}$, we think of $\mathcal{G}$ as the matching **region of** support in Radon space.[1] However, although $\mathcal{F}$ uniquely determines $\mathcal{G}$, it is **clear that** $\mathcal{G}$ uniquely determines only $\mathrm{hul}(\mathcal{F})$, not $\mathcal{F}$ itself. This is why in this paper we are **primarily** concerned with the convex support of $f(x)$, since this is what may be determined **directly** from knowledge of $\mathcal{G}$, which may in turn be estimated directly from the projections, **as** discussed in Section III.

III. REVIEW OF PREVIOUS WORK

The methods presented in [30], [31], and [32] form pieces of the hierarchical **algorithm** presented in Section IV. For completeness, in this section we provide a brief **review the** relevant results from of these papers.

A. Sinogram Estimation [30]

In [30] we developed an algorithm to estimate a full sinogram from the available **noisy** (and possibly limited-angle or sparse-angle) measured projections, assuming that **we have** prior knowledge of the convex support of the object and of the mass and center **of mass of** the function $f(\mathbf{x})$. Having this knowledge and performing appropriate **normalization and** coordinate translation we can assume that the object is centered at the origin and **has unit** mass. The estimated sinogram is taken as the solution to the variational problem, **referred** to as (V), in which we wish to choose a function $g(t,\theta)$ to minimize

$$I = \iint_{\mathcal{Y}_O} \frac{1}{2\sigma^2}(y-g)^2 \, dt \, d\theta + \iint_{\bar{\mathcal{G}}} \kappa g^2 \, dt \, d\theta + \iint_{\mathcal{Y}_T} \left[\beta \left(\frac{\partial g}{\partial t}\right)^2 + \gamma \left(\frac{\partial g}{\partial \theta}\right)^2 \right] dt \, d\theta, \qquad (13)$$

---

[1]If the support of $f(\mathbf{x})$ is not a connected set it is possible for sinogram values within $\mathcal{G}$ to be zero. Therefore, $\mathcal{G}$ it not necessarily the actual support of $g(t,\theta)$, but it does contain all the points (t,θ) for which $g(t,\theta)$ is non-zero.

subject to the equality constraints

$$J_1 \;\; = 1 = \;\; \int_{-\mathrm{T}}^{\mathrm{T}} g(t,\theta)\,dt\,, \tag{14}$$

$$J_2 \;\; = 0 = \;\; \int_{-\mathrm{T}}^{\mathrm{T}} t g(t,\theta)\,dt\,,$$

and boundary conditions

$$g(\mathrm{T},\theta) \;\; = \;\; g(-\mathrm{T},\theta) = 0\,, \tag{15}$$

$$g(t,0) \;\; = \;\; g(-t,\pi)\,,$$

where κ, β, and γ are positive constants. The numerical solution of this problem was shown in [30] to be equivalent to the maximum *a posteriori* solution of a probabilistic formulation in which the sinogram is modeled as a certain Markov random field.

The first term in I, which integrates over the set $\mathcal{Y}_O \subset \mathcal{Y}_T$ in which observed projections exist, represents a penalty that seeks to keep the estimate close to the observations. The second term integrates over the complement of the estimated region of support $\bar{\mathcal{G}}$ to attempt to keep sinogram values outside the region of support small. The final integral contains two terms involving the square of the two partial derivatives of g, which provides a smoothing effect in both the t and θ directions. The two integral constraints in (14) are exactly the mass and center of mass constraints for normalized and centered projections. The boundary conditions, stated in (15), indicate that line integrals are expected to be zero outside a disk of radius T centered at the origin, and that the sinogram is periodic, as prescribed in condition (b) of the Theorem 1.

A necessary condition for $g(t,\theta)$ to be a solution to (V) is that it satisfy the following second order partial differential equation (PDE) [36, 30]

$$\left(2\kappa\bar{\chi}_G + \frac{1}{\sigma^2}\chi_Y\right)g - 2\beta\frac{\partial^2 g}{\partial t^2} - 2\gamma\frac{\partial^2 g}{\partial\theta^2} = \frac{1}{\sigma^2}\chi_Y y - \lambda_1(\theta) - \lambda_2(\theta)t \tag{16}$$

and the additional boundary condition

$$\frac{\partial g(t,0)}{\partial t} = \frac{\partial g(-t,\pi)}{\partial t}\,, \tag{17}$$

where the two indicator functions are given by

$$\bar{\chi}_G(t,\theta) \;\; = \;\; \begin{cases} 1 & (t,\theta) \in \bar{\mathcal{G}} \\ 0 & \text{otherwise} \end{cases}\,, \tag{18}$$

$$\chi_Y(t,\theta) \;\; = \;\; \begin{cases} 1 & (t,\theta) \in \mathcal{Y}_O \\ 0 & \text{otherwise} \end{cases}\,. \tag{19}$$

In addition, $g(t,\theta)$ must still satisfy the original constraints and boundary conditions.

Equation (16) contains three unknown functions: $g(t,\theta)$, and two Lagrange multiplier functions $\lambda_1(\theta)$ and $\lambda_2(\theta)$, one for each constraint. In principle, if the Lagrange multiplier functions were known, the smoothed sinogram would simply be the solution of the partial differential equation (16), which may be solved by any of several well-known numerical methods. Good initial estimates of these Lagrange multipliers often exist [36], however, in many cases the resulting solution produces a sinogram that doesn't satisfy the constraints. In these cases, the Lagrange multiplier functions are in error and must be updated; we use the following formulas

$$\lambda_1^{k+1}(\theta) = \lambda_1^k(\theta) + \alpha\left(1 - \int_{-T}^{T} g(t,\theta)\,dt\right), \tag{20}$$

$$\lambda_2^{k+1}(\theta) = \lambda_2^k(\theta) + \alpha\left(0 - \int_{-T}^{T} tg(t,\theta)\,dt\right), \tag{21}$$

where α is a positive constant, and k is an iteration counter.

The constant α in (20) and (21) is chosen large enough so that convergence to the correct Lagrange multipliers (and, hence, the correct solution to (V)) are found quickly, yet not so large that the sequence will not converge. In our experiments, we set $\alpha = 2.0$ initially, and after each iteration we evaluate the maximum absolute error over all projections of both the mass and center of mass. If either error is greater than 0.05 and has increased since the last iteration then we divide α by 2 prior to the next iteration, but do not allow α to become less than 0.05. This empirical adaptive schedule has yielded a good rate of convergence for the problems we have studied. Bertsekas [37] describes more precisely the trade-offs in the selection of α, and relates this generic primal-dual method to the method of multipliers, about which a great deal of theory is known.

B.   Support Vector Estimation [31, 32]

Even in the full-data problem, only a finite number n_v of projections are measured; therefore, we can only hope to measure a finite number of support values from these projections. This corresponds to measuring a sampled version of the support function of the convex support of the object [see Equation (11)]. We consider a finite number $M = 2n_v$ of angles $\theta_i = 2\pi(i-1)/M$, $i = 1,\ldots,M$, spaced uniformly over $[0, 2\pi)$. A *support vector* h is a vector made by organizing the values of a support function $h(\theta)$ sampled at these angles, yielding

$$\mathbf{h} = [h(\theta_1) \quad h(\theta_2) \quad \ldots \quad h(\theta_M)]^T . \tag{22}$$

Not all vectors in $\mathbb{R}^M$ are support vectors, however. In particular, in [31] we proved that a vector $h \in \mathbb{R}^M$ ($M \geq 5$) is a support vector if and only if

$$h^T C \leq [0 \ldots 0], \tag{23}$$

where C is an M by M matrix given by

$$C = \begin{bmatrix} 1 & -k & 0 & & & -k \\ -k & 1 & -k & \ldots & \ldots & 0 \\ 0 & -k & 1 & & & \vdots \\ \vdots & 0 & -k & \ldots & \ldots & 0 \\ 0 & \vdots & & & & -k \\ -k & 0 & 0 & & & 1 \end{bmatrix} \tag{24}$$

and $k = 1/(2\cos(2\pi/M))$.

As one would expect, if we only have available uncertain measurements of support values (obtained, for example, using the knot-location algorithm developed in Section IV), it is possible that the resulting measured support vector will be inconsistent with any object in the plane. The support vector estimation principles and methods described in [31] and [32] and reviewed in this section provide a framework for consistent estimation of convex sets, based on noisy support vector measurements, and yield methods that are used here in the hierarchical algorithm. The formulations in [31], [32] also allow us to consider problems in which we wish to reconstruct the M-dimensional support vector h based on $K \leq M$ noisy measurements, modeled as

$$\tilde{z} = Sh + v, \tag{25}$$

where v is a zero-mean jointly Gaussian vector with covariance $\text{diag}[(\sigma_e^2)_j]$ and S is a $K \times M$ "selection" matrix specifying the components of h that are measured (thus allowing us to consider limited-angle and sparse-angle cases). The procedure we describe in Section IV not only produces $\tilde{z}$ but also the variances of the components of v.

The log likelihood of h, given the above observation model, is given by

$$l(h) = -\frac{1}{2\sigma^2}(\tilde{z} - Sh)^T(\tilde{z} - Sh) - \frac{1}{2}\ln\left|2\pi\sigma^2 I_k\right| \tag{26}$$

which may be written as

$$l(h) = \frac{1}{2\sigma^2}(z - h)^T D(z - h) - \frac{1}{2}\ln\left|2\pi\sigma^2 I_k\right| \tag{27}$$

where

$$D = S^T S, \text{ and} \tag{28}$$

$$z = S^T \tilde{z}. \tag{29}$$

If we have a full set of observations, i.e. if $S = I$, we can produce the maximum likelihood (ML) support vector estimate by finding h which maximizes the log-likelihood subject to the support vector constraint of (23). We call this the *constrained ML* support vector estimate. Alternatively, we may include prior geometric knowledge — and *must* do so if the observations are not complete (i.e., $S \neq I$) — by specifying *a priori* probabilities on the set of all feasible support vectors or by assuming the object to be a known geometric shape with unknown parameters. The maximum *a posteriori* (MAP) solution is then given by

$$h_{MAP} = \underset{h \mid h^T C \leq 0}{\text{argmax}} \quad -\frac{1}{2\sigma^2}(z - h)^T D(z - h) + \ln p(h), \tag{30}$$

where we have dropped the constant term in (27) and $p(h)$ is a prior probability defined for feasible support vectors. If the logarithm of the prior probability is linear or quadratic then (30) is a quadratic program (QP).

There is a rich set of possible probabilistic and deterministic specifications for prior shape information. For example, in [32] we considered the possibility that the object is known to be nearly a circle or an ellipse, with both known and unknown ellipse parameters. We defined several prior probabilities called the *scale invariant* (SI) priors, which capture this prior shape information in different ways. The SI maximum area (SIMA) prior probability favors objects whose area is largest given a fixed circumference, thus giving nearly circular objects larger *a priori* probability. The SI closest (SIC) prior probability favors objects whose support vectors are near (in the Euclidean sense) to the support vector of a circle, thus also favoring nearly circular objects, but in a different way. The SI close-min (SICM) prior probability favors objects whose boundaries are smooth in the sense of having the maximum minimum discrete radius of curvature given a fixed circumference. Again, circles have the highest prior probability in the SICM prior. Finally, the ellipse-based SI closest (ESIC) prior probability favors objects whose support vectors are near to a particular ellipse, whose eccentricity and orientation must be specified in advance.

A particularly interesting alternative approach to prior shape specification, called the *joint ellipse* (JE) algorithm, was also developed in [32]. Here, the support vector of the object is assumed to be near that of an ellipse but the particular ellipse is not known. Therefore, the objective of the support vector estimation algorithm is to jointly estimate an

ellipse and a support vector from the data. Surprisingly, this approach performed nearly as well in simulations as the ESIC algorithm, which assumed knowledge of the true ellipse.

Estimating the support vector from incomplete data requires prior knowledge; hence, in the limited-angle and sparse-angle cases considered in the experiments of Section V, we specify one of SI approaches or the JE approach. Each approach also requires the specification of an additional parameter τ, which may be thought of as a regularization parameter, which gives the strength of imposition of prior knowledge. In this work we choose these parameters empirically, but note that it is often possible to estimate regularization parameters directly from the data, using for example the method of cross validation [38].

IV.   HIERARCHICAL ALGORITHM

In this section, we present a synthesis of the sinogram and convex support estimation methods reviewed above with new mass, center of mass, and support value estimation methods described below in order to produce a complete, hierarchical reconstruction algorithm. We assume that the observations are given by

$$y(t_i, \theta_j) = g(t_i, \theta_j) + n(t_i, \theta_j) \tag{31}$$

where the indices $i \in \{1, \ldots, n_d\}$ and $j \in \mathcal{J} \subset \{1, \ldots, n_v\}$ index points on the regular rectangular lattice in the domain $\mathcal{Y}_T$. The set $\mathcal{J}$ contains indices corresponding to the angular positions of the observed projections; we use J to denote the number of such projections. The noise samples $n(t_i, \theta_j)$ are zero-mean white Gaussian random variables, independent between lattice sites.

A block diagram of the new algorithm is shown in Figure 5. Each block represents a significant stage which either estimates a new parameter or set of parameters or transforms the data in some fashion; blocks that correspond to prior work are shaded. Where required, an estimate of the reliability of the information is also passed between the blocks. In this way, poor estimates are not viewed as perfect by subsequent processing stages, and extremely good estimates are given greater weight in subsequent blocks.

We now give an overview of the processing sequence of the hierarchical algorithm, each new element of which is described in detail in subsequent sections. As shown in Figure 5, the overall processing is divided into four stages: (A) Sinogram Preconditioning, (B) Sinogram Restoration, (C) Sinogram Postconditioning, and (D) Object Reconstruction. In (A) the mass and center of mass are estimated directly from the available measured projections y, using methods described in a subsequent section. These quantities are used to center the

coordinate system by shifting each projection to correspond to an object centered at the origin and to normalize each of the projections to unit mass, yielding $\tilde{y}$. Stage (B) forms the bulk of the processing, with support vector estimation and MAP sinogram estimation — two methods reported in previous work — and two new steps required to support these procedures. Here, the the block labeled *Support Vector Estimation* estimates a support vector $\hat{h}$ as described in Section III.B, which is used to define a corresponding segmentation of the sinogram into $\mathcal{G}$ and $\bar{\mathcal{G}}$. These sets are used by the block labeled *MAP Sinogram Estimation* as described in Section III.A to estimate a full sinogram. Support value estimates, required for support vector estimation are provided by the new block labeled *Support Value Estimation*, which is described in a subsequent section. This procedure requires a threshold which is adaptively estimated in the block labeled *Threshold Estimation*, also described below. Stage (C) takes the restored sinogram as input and rescales its mass to the estimated mass and restores the coordinate system to the estimated center of mass. Finally, Stage (D) performs convolution backprojection on the restored sinogram to reconstruct an object estimate $\hat{f}$.

As shown in Figure 5, the hierarchical algorithm requires several user inputs: γ, β, κ, used by the MAP sinogram estimation block, and τ, prior shape information used by the support vector estimation block. The parameters γ and β specify *a priori* sinogram horizontal and vertical smoothness, respectively; κ specifies a measure of confidence in the support vector estimate. The prior shape information τ for support vector estimation includes the choice of a method and one or more parameters associated with the chosen method. The overall performance of the hierarchical algorithm is affected by how well the parameters we use represent the truth; in our experiments they are empirically adjusted to match the class of objects and imaging geometry used in the experiments. (See [31] and [32] for detailed discussions of the effects of parameter selection on the MAP sinogram estimation algorithm and the support vector estimation algorithm, respectively.)

A. Mass and Center of Mass Estimation

The mass and center of mass of the object are estimated directly from the projections using least squares. We assume that the mass constraint given by (4) may be approximated by the summation

$$\frac{2T}{n_d} \sum_{i=1}^{n_d} g(t_i, \theta_j) = m \quad \forall j \, . \tag{32}$$

Substituting the observed projections $y(t_i, \theta_j)$ for the true projections $g(t_i, \theta_j)$ yields a system of equations that may be solved for m using least squares. Accordingly, the mass estimate

is

$$\hat{m} = \frac{2T}{n_d J} \sum_{j \in \mathcal{J}} \sum_{i=1}^{n_d} y(t_i, \theta_j), \tag{33}$$

which is proportional to the average of all of the observed line integrals.

To estimate the center of mass, we first approximate the integral in (5) by the summation

$$c_j = \frac{1}{m} \frac{2T}{n_d} \sum_{i=1}^{n_d} t_i g(t_i, \theta_j). \tag{34}$$

Substituting the mass estimate $\hat{m}$ for m, the observed projections $y(t_i, \theta_j)$ for $g(t_i, \theta_j)$, and $\mathbf{c} \cdot \omega_j$ for c_j yields a system of linear equations with unknown $\mathbf{c}$, the center of mass of the object. This system may also be solved using least squares, yielding the center of mass estimate

$$\hat{\mathbf{c}} = (\mathbf{A}^\mathrm{T} \mathbf{A})^{-1} \mathbf{A}^\mathrm{T} \mathbf{b} \tag{35}$$

where

$$\mathbf{A} = \begin{bmatrix} \cos \theta_1 & \sin \theta_1 \\ \vdots & \vdots \\ \cos \theta_J & \sin \theta_J \end{bmatrix} \tag{36}$$

and

$$\mathbf{b} = \begin{bmatrix} \tilde{c}_1 \\ \vdots \\ \tilde{c}_J \end{bmatrix}, \tag{37}$$

where

$$\tilde{c}_j = \frac{1}{\hat{m}} \frac{2T}{n_d} \sum_{i=1}^{n_d} t_i y(t_i, \theta_j). \tag{38}$$

B. Object Centering and Mass Normalization

Using the mass and center of mass estimates, the measured projections are shifted and scaled so that the center of mass of the object is at the origin and the mass is unity. The modified projections are given by

$$\tilde{y}(t, \theta) = \frac{1}{\hat{m}} y(t - \hat{\mathbf{c}} \cdot \omega, \theta). \tag{39}$$

This processing is done so that the convex support estimation stage and sinogram estimation stage may assume the object to have unit mass and to be centered at the origin. After the sinogram processing in Stage (B) of the hierarchical algorithm, the full set of estimated projections are shifted back using a similar equation.

C. Support Value Estimation

The convex support estimate is made in two stages. The first stage measures two support values for each available projection (estimates of t_- and t_+ in Figure 3 for each available θ). Since the projections are noisy, these measurements will be inaccurate, and therefore must be considered to be noisy measurements of the true support values of the convex support of the object. The second stage estimates a complete feasible support vector from these support value measurements using the methods reviewed in Section III.B. In this section we present an algorithm for support value measurement based on detection of jumps in linear systems. A method to adaptively estimate the error in these measurements is also presented, and its use in consistent support vector estimation is discussed.

1) Knot-Location Method

In this section we model a projection $g(t)$ as a continuous piecewise-linear waveform as shown in Figure 6. Such a function is called a linear *spline*; it is composed of a set of linear functions that connect a series of points called *knots*, so that the resultant function is continuous but its slope has an abrupt change at each of the knots. Two support values for this projection are determined by estimating the positions, t_- and t_+, of the two outer knots.

Our *knot-location method* is based on the generalized likelihood ratio (GLR) techniques developed by Willsky and Jones [39] for detecting abrupt changes in dynamic systems, and later applied to spline estimation by Mier-Muth and Willsky [40]. The approach is to run a Kalman filter, starting at $t = -T$, assuming an underlying signal model corresponding to a *linear ramp* waveform (initialized with with slope=0). The innovations of the Kalman filter, which should be a zero-mean white Gaussian sequence, provide the basis for estimating the point at which the true signal deviates from the assumed model — i.e. the point at which the slope changes suddenly. The first such point yields the point t_-. The point t_+ is found by running the filter backwards starting at $t = T$ and finding the first knot in the backwards direction.

As discussed in [40], a linear spline with a single knot can be described by the following two-dimensional discrete state equation

$$\mathbf{x}(i+1) = \Phi\mathbf{x}(i) + \alpha\delta(i+1-k)\mathbf{f} \tag{40}$$

where $\delta(\cdot)$ is the discrete impulse function, α is the height of the discontinuity, k is the

discrete position of the knot $(1 \leq k \leq n_d)$, and

$$\Phi = \begin{bmatrix} 1 & 2T/n_d \\ 0 & 1 \end{bmatrix} \tag{41}$$

$$f = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \tag{42}$$

The problem of fitting such a linear spline to an observed data sequence $y(i)$ corresponds to estimating the parameters of (40) — namely the initial conditions and slope, and the knot location k and jump height α — assuming that the data are noisy measurement of the spline, i.e.

$$y(i) = h^T x(i) + v(i), \tag{43}$$

where $h^T = [\, 1 \quad 0 \,]$, and $v(i)$ are zero-mean white jointly Gaussian random variables with variance $R = \sigma^2$. In our problem the $y(i)$ represent the projection measurements $y(t_i, \theta_j)$ as a function of i for each fixed θ_j.

The first step in the knot location algorithm is to run the following Kalman filter on the data:

$$\hat{x}(i|i-1) = \Phi \hat{x}(i-1|i-1), \tag{44}$$

$$\hat{x}(i|i) = \hat{x}(i|i-1) + K(i)\gamma(i), \tag{45}$$

$$\gamma(i) = y(i) - h^T \hat{x}(i|i-1), \tag{46}$$

where $\hat{x}(i|l)$ is the best estimate of $x(i)$ given $y(1), \ldots, y(l)$, $\gamma(i)$ is the innovations sequence, and $K(i)$ is the Kalman filter gain. Assuming that there is no jump (slope discontinuity), the innovations sequence is a zero-mean, white, jointly Gaussian random sequence whose variance $V(i)$ is given by

$$V(i) = h^T P(i|i-1)h + R, \tag{47}$$

where the error covariance $P(i|i-1)$ may be computed together with the Kalman gain $K(i)$ using the following recursive algorithm:

$$P(i|i) = [I - K(i)h^T]P(i|i-1), \tag{48}$$

$$K(i) = P(i|i-1)hV^{-1}(i), \tag{49}$$

$$P(i+1|i) = \Phi P(i|i)\Phi^T. \tag{50}$$

Because a projection is zero outside the disk of radius T and since we may take T to be as large as necessary, we know with certainty the initial spline parameters and thus initialize

the Kalman filter as follows:

$$\hat{\mathbf{x}}(0|0) \;=\; \hat{\mathbf{x}}(1|0) = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \tag{51}$$

$$\mathbf{P}(0|0) \;=\; \mathbf{P}(1|0) = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}. \tag{52}$$

In order to determine whether a jump has occurred we examine the innovations sequence which will deviate from the statistics given above if a jump takes place. In particular, given our single jump model of (40) it can be shown that the true innovations sequence takes the form [39]

$$\gamma(i) = \mathbf{G}(i,k)\mathbf{f}\alpha + \tilde{\gamma}(i)\,, \tag{53}$$

where $\tilde{\gamma}(i)$ is the zero-mean, white, and Gaussian innovations if there is no jump and $\mathbf{G}(i,k)$ is the *jump signature matrix* given by

$$\mathbf{G}(i,k) = \mathbf{h}^{\mathrm{T}} \left[\Phi^{i-k} - \Phi\mathbf{F}(i-1,k) \right]\,, \tag{54}$$

$$\mathbf{F}(i,k) = \mathbf{K}(i)\mathbf{G}(i,k) + \Phi\mathbf{F}(i-1,k)\,, \tag{55}$$

where $\mathbf{G}(i,k)$ and $\mathbf{F}(i,k)$ are both 0 for $i < k$ and $\mathbf{F}(i,i) = \mathbf{K}(i)\mathbf{h}^{\mathrm{T}}$.

Equation (53) is the key to the GLR knot-location method. Through this equation, we see how to form the ML estimate of α assuming a jump occurred at time k in the filter's past for each current time index i of the Kalman filter. Actually, to reduce the required computation, at each point i we look for possible jumps only over a trailing window $W(i) = \{i-1, i-2, \ldots, i-N\}$ of length N in the filter's past. Then using the ML estimate of α for each $k \in W(i)$, we form the GLR for the hypothesis that a jump actually occurred at k. If the GLR exceeds a preset threshold then a jump is deemed to have occurred. The above calculations are given by the following equations each evaluated for all $k \in W(i)$ [40]

$$\mathbf{C}(i,k) \;=\; \sum_{j=i}^{i} \mathbf{G}^{\mathrm{T}}(j,k)V^{-1}(j)\mathbf{G}(j,k)\,, \tag{56}$$

$$\mathbf{d}(i,k) \;=\; \sum_{j=k}^{i} \mathbf{G}^{\mathrm{T}}(j,k)V^{-1}(j)\gamma(j)\,, \tag{57}$$

$$\hat{\alpha}(i,k) \;=\; \frac{\mathbf{f}^{\mathrm{T}}\mathbf{d}(i,k)}{\mathbf{f}^{\mathrm{T}}\mathbf{C}(i,k)\mathbf{f}}\,, \tag{58}$$

$$l(i,k) \;=\; \frac{(\mathbf{f}^{\mathrm{T}}\mathbf{d}(i,k))^2}{\mathbf{f}^{\mathrm{T}}\mathbf{C}(i,k)\mathbf{f}}\,, \tag{59}$$

where $\hat{\alpha}(i,k)$ is the ML estimate of α assuming that a jump occurred at time k, and $l(i,k)$ is the logarithm of the generalized likelihood ratio for this event. The best estimate of the

location of a jump is then given by

$$\hat{k}(i) = \underset{k \in W(i)}{\mathrm{argmax}}\, l(i, k).$$

(60)

Then, to decide whether a jump has actually taken place we use the following threshold rule

$$l(i, \hat{k}(i)) \begin{array}{l} > \varepsilon \quad \text{Jump} \\ \leq \varepsilon \quad \text{No Jump} \end{array}.$$

(61)

2) Selection of Jump Threshold

Specification of the GLR threshold ε is an important consideration since if too low, noise will often produce an inaccurate knot estimate, and if too high, it is possible that no knot will be found in the entire projection. We now describe an adaptive method to choose this threshold that has worked well in practice.

Since the mass of each projection is the same, we might expect that a projection with a small support width would rise rapidly at the support values in order to include the required mass, and this would correspond to a large value of α and a corresponding large value of $l(i, \hat{k}(i))$. In contrast, a projection with a larger support width might rise less rapidly, and would correspond to a smaller value of $l(i, \hat{k}(i))$. A simple estimate of the width of a projection is given by the approximate second moment of a normalized (to unit mass), shift-corrected projection as follows

$$m_2(\theta_j) = \sum_{i=1}^{n_d} t_i^2 \max\{0, \tilde{y}(t_i, \theta_j)\}.$$

(62)

The max{ } function is included since it possible that elements of $\tilde{y}$ are negative, and therefore that $m_2(\theta)$ might otherwise be negative.

The second moment in (62) is roughly equivalent to a variance calculation, and the quantity

$$p(\theta) = \sqrt{m_2(\theta)}$$

(63)

is analogous to a standard deviation, which serves as an approximate measure of the width of the projection. If p is large, then the projection is wide and the slope change at the support value is probably small. Therefore, we want to specify a GLR threshold ε that is relatively small. Using similar reasoning we conclude that for small p, the threshold ε should be large. After some experimentation we have chosen the function $\varepsilon(p)$ depicted in Figure 7. Since p is a measure of the width of the entire projection, ε is used as the threshold value for both the forward and backward stages of the knot-location algorithm (i.e. for measuring both t_- and t_+).

3) Performance of the Support Value Estimation Algorithm

It is important in our hierarchical approach to be able to assess the performance of the support value measurement method so that this information may be used by the subsequent support vector estimation stage. The accuracy of the support value measurements clearly depends not only on the variance of the additive noise but also on the characteristics of the underlying projection, particularly at or around the true support value. Indeed if these characteristics were summarized, say in a template model, one could compute a Cramer-Rao bound specifying these dependencies in a quantitative manner. Unfortunately, these characteristics can vary widely from projection to projection even for the same object. Thus it is essential that we have a method for determining the quality of our support measurements directly from the projection data. In this section we present a method to obtain such an estimate of the error variance arising from the knot-location algorithm described above. In doing this we do not assume any prior shape information at this stage, so our estimate is made completely on the basis of statistics available during processing, and, in particular, on the shape of the log-likelihood function. If the log-likelihood function is sharply peaked at its maximum then we presume that it is a good estimate; if it has a shallow maximum then we presume that the estimate is not as good. Following this principle, we fit a downturned quadratic centered at $\hat{k}$ to the log-likelihood function that was evaluated over the window $W(i)$, and the coefficient of the quadratic term yields our error variance estimate. It is worth noting that this quadratic fitting method can be viewed as a signal-adaptive estimate of the Cramer-Rao bound.

Let $\hat{k} \in W(i)$ be our estimate of the knot location, made when the Kalman filter has progressed to the ith index. We wish to fit a downturned quadratic of the form

$$\hat{l}(k) = -a(k - \hat{k})^2 + c \tag{64}$$

to the data $l(i, k)$ so that, in particular, we may determine a. To make this fit, we minimize

$$\sum_{k=1}^{N} (\hat{l}(k) - l(i, k))^2 = \sum_{k=1}^{N} \left(-a(k - \hat{k})^2 + c - l(i, k) \right)^2 ,$$

with respect to a and c, yielding

$$a = \frac{\sum_{k=1}^{N} (c - l(i, k))(k - \hat{k})^2}{\sum_{k=1}^{N} (k - \hat{k})^4} , \tag{65}$$

$$c = l(i, \hat{k}) . \tag{66}$$

Our estimate of error variance σ_e^2 for the support estimate is then given by

$$\sigma_e^2 = \frac{1}{2|a|} . \tag{67}$$

D. Results of Support Value Estimation

Figures 8 and 9 show simulation results in which the knot-location algorithm is used to estimate support values for several noisy projections. Both figures show noisy projections of an ellipse centered at the origin, with major semiaxis radius of 0.806 and minor semiaxis radius of 0.242 — these are the same dimensions as the M I T ellipse, shown in **Figure 2**. Figure 8 shows the narrowest projection — the support values are -0.242 and 0.242 — with different noise variances. Figure 9 shows the widest projection — the support values are -0.806 and 0.806 — also with different noise variances. The positions of the true support values are indicated by the vertical dotted lines in each of the panels. The indicated noise standard deviation σ in each of the panels of Figures 12 and 13 corresponds to that of noise added to *full* sinograms to achieve the SNR of (a) 100.0dB, (b) 10.0dB, (c) 3.0dB, and (d) 0.0dB (see Section V). The same underlying unit variance noise sequence was used for all four projections in each figure, which accounts for the similarity in the noise structure.

Support value estimates are indicated by diamond markers in each of the panels in Figures 8 and 9. The error bar centered directly above each of these symbols has a length of two standard deviations; in several cases the error bars are very short and appear merely as points. In most cases, the estimated values are within three standard deviations of the true values; but, as one would expect, the size of the error increases as the noise variance grows. Also, the error bars are longer in Figure 9 than those in corresponding panels in Figure 8. This agrees with our intuitive reasoning that it should be more difficult to detect the onset of a projection of the broadside of the ellipse versus the head-on projection, given the same noise variance (since the broadside projection has a more gradual rise than the head-on projection).

E. Support Vector Estimation

In the block diagram of Figure 5, the support value measurements $\{\hat{k}_i \mid i = 1, \ldots, 2n_v\}$ feed into a block labeled *Support Vector Estimation*. In addition to the measurements themselves, their estimated error variances $\{(\sigma_e^2)_i \mid i = 1, \ldots, 2n_v\}$ are also passed along. The support value measurements are used as support vector observations to form vector $\tilde{z}$ [see Equations (25) and (29)]; and the estimated error variances provide the variances of the additive noise v [see Equation (25)]. Given these data, a consistent support vector $\hat{h}$ is estimated using one of the methods described in Section III.B. This support vector is then passed to the *MAP Sinogram Estimation* block (see Figure 5) which uses it to form an estimate $\hat{\mathcal{G}}$ of the sinogram support region using Equation (12).

Figure 10 shows a noisy sinogram in each panel together with the full set of knot-location support value measurements (shown using thin white curves) and the constrained ML support vector estimate (shown using thick white curves), obtained by maximizing the log likelihood given by Equation (26). The signal-to-noise ratio (SNR) of the sinogram, using the definition of SNR given in Section V, in each panel is given by (a) 100.0dB, (b) 10.0dB, (c) 3.0dB, and (d) 0.0dB. The set defined between the top and bottom thick white lines is $\hat{\mathcal{G}}$, which should ideally contain all pixels with non-zero values. As the noise level increases from Figures 10a–10d, the performance of the knot-location algorithm gets noticeably worse, but the constrained ML support vector estimate does not show the same qualitative degradation. Since the constrained ML algorithm does not use any prior geometric information, this shows that utilization of support vector consistency alone adds considerable robustness to errors in the knot-location measurement process.

Confidence in the overall determination of $\hat{\mathcal{G}}$ is specified by the parameter κ used for sinogram restoration (see Equation (13) and Figure 5). A large κ indicates great confidence in $\hat{\mathcal{G}}$; a small κ indicates little confidence in $\hat{\mathcal{G}}$. Since choosing a large κ causes data outside $\hat{\mathcal{G}}$ to be largely ignored it is most dangerous to do so when $\hat{\mathcal{G}}$ is too small — i.e., when we have estimated a convex hull for $f(x)$ that does not strictly include the true convex hull of $f(x)$. Therefore, in Section V, we evaluate two approaches for the selection of κ, one that uses a small κ ($\kappa = 5.0$) and the given estimate $\hat{\mathcal{G}}$ and one that uses a large κ ($\kappa = 10,000$) with a modified region of support $\hat{\mathcal{G}}_m$, which is selected to contain $\hat{\mathcal{G}}$.

The modified region of support $\hat{\mathcal{G}}_m$ is constructed by adding to each estimated support value some fraction of its own estimation error standard deviation. In this way, for larger estimation error variances, the boundary will be further away from the estimated support value, and its effect on the estimated sinogram values near the boundary will be reduced. Letting $\hat{h}(\theta)$ be the estimated support function and $\hat{\sigma}_e(\theta)$ be the estimation error variance of the support estimate at angle θ, we define the modified sinogram region of support to be

$$\hat{\mathcal{G}}_m = \left\{ (t, \theta) \in \mathcal{Y}_T \mid \hat{h}(\theta) + \xi \hat{\sigma}_e(\theta) \leq t \leq -\hat{h}(\theta + \pi) - \xi \hat{\sigma}_e(\theta + \pi) \right\}, \qquad (68)$$

where ξ is an arbitrary positive parameter (typically 1 or 2).

V. EXPERIMENTAL RESULTS

Figure 2a shows the 81 by 81 pixel M I T ellipse object used in the first set of experiments in this section. An 81 by 60 noise-free sinogram, derived from an approximate strip-integration of an analytic representation of the M I T ellipse, is shown in Figure 2b. This object was chosen for experimentation because the loss of data over different angular regions affects the reconstructions in different ways. For example, the absence of line integrals parallel to the long axis of the ellipse causes a lack of information related to the narrow dimension of the ellipse, but retains information about the letters inside the ellipse. In contrast, the absence of line integrals parallel to the short axis of the ellipse obscures the letters, but reveals the narrowness of the ellipse.

A. Complete Observations

To synthesize noisy observations we add independent samples of zero-mean Gaussian noise with variance σ^2 to each element of the true sinogram of Figure 2a. The resulting sinogram has signal-to-noise ratio (SNR) defined as

$$\text{SNR} = 10 \log \frac{\frac{\pi}{n_v} \frac{2T}{n_d} \sum_{j=1}^{n_v} \sum_{i=1}^{n_d} g^2(t_i, \theta_j)}{\sigma^2}, \tag{69}$$

where $g(t_i, \theta_j)$ is the true sinogram. Figure 11a shows a 3.0dB sinogram of the M I T ellipse and Figure 11b shows a reconstruction from these complete data using convolution backprojection.

Figure 12a shows the sinogram estimate obtained from the hierarchical algorithm using the full noisy 3.0dB sinogram of Figure 11a, with $\gamma = 0.05$, $\beta = 0.01$, and $\kappa = 5.0$. The estimated segmentation, made using the knot-location algorithm to estimate support values followed by the constrained ML algorithm to estimate a feasible support vector, is superposed on the figure. The reconstructed image of Figure 12b results from CBP applied to the restored sinogram in Figure 12a. Figures 12c and 12d show the sinogram estimate, segmentation, and reconstruction obtained using the same data and parameters, except that $\kappa = 10,000$ and the modified segmentation $\hat{\mathcal{G}}_m$ was used with ξ equal to one standard deviation.

The reconstructions shown in Figure 12, panels (b) and (d), show three major differences — although even these differences are subtle in these examples. First, the object in (d) has greater contrast between the object and its background than in (b); this primarily reflects the fact that κ is much larger in (d) than in (b). Second, there is less contrast in the *interior* of

the ellipse in panel (d) than in (b). Third, there is a slight two-tiered effect on certain parts of the boundary of the object in (d), where the effects of two boundaries — the segmentation boundary and the true boundary — are apparent. All three of these effects were observed in [30], and may be viewed as fundamental properties of this sinogram restoration algorithm. In particular, when κ is very large (e.g. 10,000) one can expect an increase in overall contrast between the object and its background, but this is at the expense of loss of contrast in the interior (due, in part, to the hard, constant-mass constraint — see [30]).

From this experiment, it appears that as long as κ is small enough (e.g. 5.0), there is some increase in overall contrast and little of the two-tiered effects, or loss of interior contrast. Therefore, in the experiments presented in the following section — which include limited-angle and sparse-angle cases as well as results for a different object — we use unmodified support value estimates and $\kappa = 5.0$.

B. Limited-angle and sparse-angle observations

Two limited-angle and two sparse-angle cases are considered, both using selected projections from a 10dB noisy MIT sinogram which is not shown. One limited-angle case observes the left-most 40 (out of 60) projections and the other observes the right-most 40 projections. This is generally considered to be a severe limited-angle problem where even missing as little as 1% of the data can produce severe artifact [14]. One sparse-angle case observes every 4th projection starting with the first column, for a total of 15 projections. The second sparse-angle case observes every 6th projection for a total of 10 projections. Reconstructions using convolution backprojection — with all unobserved projections set to zero — for these four cases are shown in Figure 13. Figures 14 and 15 show the results of applying the hierarchical algorithm to these two limited-angle and two sparse-angle data sets. The panels correspond to the those in Figure 13. Figure 14 shows restored sinograms with estimated segmentations superposed and Figure 15 shows the objects reconstructed (using CBP) from the restored sinograms.

Although fewer projections are observed than are actually present in the restored sinogram, the support estimation procedure obtains estimates for *all* of the projections. The estimates obtained for the missing projections — these may be thought of as *interpolated* support values — strongly rely on the prior knowledge provided by the geometric information assumed as prior knowledge. In these cases, the Scale-Invariant Maximum Area prior is used (see Section III.B and [32]). When we do not observe the narrow view of the ellipse as in Figures 13a, 14a, and 15a the support estimate is not as good because we assumed that

the objects are likely to be circular. This knowledge, reflected in the superposed sinogram segmentation is more easily seen in the reconstructed object of Figure 15a, where the ellipse bulges outward away from the long axis.

The results of these simulations should be compared to the (unprocessed) CBP reconstructions shown in Figure 13. We see that the results presented here are dramatically better in each case.

C.   Two-disk object.

In the following set of simulation studies we introduce a new object which consists of two isolated disks rotated off the vertical axis by 45.0 degrees, as shown in Figure 16a. Figure 16b shows a noise-free sinogram of this object. In several ways, this figure is similar to the M I T ellipse: it is oriented in a similar manner and it may be viewed broadside (where in this case it allows us to see that there are two distinct objects — although we do not explicitly use this fact in our processing) and head-on.

Figures 17a and 17b show two reconstructions for the limited-angle case in which the left-most 40 projections of Figure 16b are observed in noise (SNR=10.0dB). Figure 17a is the CBP reconstruction and Figure 17b is the hierarchical reconstruction using the identical parameters (and support vector prior) as in the above limited-angle experiments. Figures 17c and 17d show two reconstructions for the limited-angle case in which the right-most 40 projections are observed. As above, Figure 17c results from CBP and Figure 17d results from the hierarchical algorithm. The convex support estimate is superposed on the reconstructions in Figures 17b and 17d.

As in the limited-angle studies with the M I T ellipse, the performance of the algorithm differs substantially between the two different limited-angle cases. In particular, where the majority of the views are obtained from the left side of the sinogram (top two panels in Figure 17), the narrow dimension of the two disks is not observed, and the resultant support estimate is too wide. Once again, this reflects the fact that the most likely shapes (using the SI Maximum Area prior) are circular. The consequence of this property is to produce a narrow band (or swath) of bright values in the reconstruction, which follows just inside the estimated boundary (see Figure 17b). It is an artifact that can just barely be discerned in the direct CBP reconstruction shown in Figure 17a, so that it definitely results from our enhancement efforts. The reason for this bright band lies primarily in the mass constraint, which together with the support constraint, causes most of the mass to be in the region of estimated support, and also because of the horizontal smoothing coefficient, which causes

rotational swirling. The brightness of the band, however, is substantially lower than that of the two disks, and in fact its appearance is heightened by the superposed support boundary.

The hierarchical reconstruction in Figure 17d, corresponding to measurements of the right-most 40 projection, shows significant improvement over the CBP estimate shown in Figure 17c. It also shows an excellent convex support estimate, which largely reflects the fact that the missing projections correspond to the broadside views, and their support values have a circular curvature at this point. The only significant artifact that is present in this reconstruction — notice that the banding artifact is *not* present — is a very slight tendency for the disks to be elongated towards each other in the center. This arises from vertical smoothing of the sinogram and the fact that many of the projections from the broadside views of the two disks are missing, and therefore had to be interpolated.

D. Ellipse-based Priors.

We have now seen two examples of limited-angle observations in which the narrow views were missing, and this resulted in poor convex support estimation and correspondingly poor reconstructions. In our final simulations we use the prior geometric information that the convex support is nearly that of an ellipse to improve the convex support estimation and hence the reconstructions. Figure 18 shows several results in which the Ellipse-Based support vector estimation of [32] is used instead of the SI Maximum Area algorithm to estimate the convex support from the available support value estimates. In this figure, the top two panels are reconstructions of the M I T ellipse from the left 40 projections of a 10.0dB SNR noisy sinogram. The bottom two figures are reconstructions of the two-disk object given noisy observations (SNR=10.0dB) of the left 40 projections of Figure 16b. Figures 18a and 18c use the joint ellipse (JE) support vector estimation algorithm with $\alpha = 0.5$, while Figures 18b and 18c use the ellipse-based support vector estimation algorithm (ESIC) with $\bar{\varepsilon} = 0.9$ and $\bar{\phi} = -45.0^\circ$ [32]. The critical difference between these two approaches is that while both methods assume that the true support is nearly that of an ellipse, the JE method requires no particular information about the ellipse, while ESIC requires both the orientation and eccentricity. Note that the true value of $\bar{\varepsilon}$ for the M I T ellipse is 0.95, so that ESIC is given a smaller eccentricity than the truth in this case. The support estimates are superposed over each reconstruction using white curves.

To evaluate these results, Figures 18a and 18b should be compared to Figure 15a; and Figures 18c and 18d should be compared to Figure 17b. Figures 18a and 18b are obvious improvements over Figure 15a, with 18a showing a remarkable result in light of the fact that

the only prior information given is that the object is shaped like an ellipse. It is fairly clear also that Figures 18c and 18d are better than Figure 17b, despite the fact that the two-ball object is not elliptical. (Note: Figures 18c and 18d should not be compared to Figure 17d, since Figure 17d results from a different set of observations.)

It should be noted that the difference between these figures we have compared above is the nature and specificity of the prior geometric knowledge. In particular, in these ellipse-based experiments the width of the support estimates are much narrower than those obtained using the SI Maximum Area algorithm in the previous examples, and this has the effect of concentrating most of the reconstructed energy within a region that more closely approximates the true region of convex support. Also, since the available views of the M I T ellipse and the two-disk object allow us to see the details of, respectively, the internal letters and the gap between the two objects, the clarity of these features remains good in these reconstructions. This demonstration is clear example of how support information — estimated in a hierarchical fashion — can lead to substantially improved reconstructions over those obtained by conventional methods.

VI. DISCUSSION

We have demonstrated a method to estimate and hierarchically incorporate geometric information in a reconstruction algorithm designed for noisy and limited-angle or sparse-angle tomography. The method is based on estimation principles, incorporating prior probabilistic information and consistency conditions to overcome problems resulting from insufficient data.

Many variations of the basic algorithm may be considered, and several are discussed in [36]. One variation is to incorporate more specific information about the shape of the object's convex hull. For example, with the additional knowledge that the object is an ellipse, but without knowing the size, orientation, or position or the ellipse, the algorithm produces nearly perfect support estimation in the experimental geometries and SNR used in Section V. Another variation is to incorporate more than just the two consistency conditions given by the mass and center of mass constraints. One approach, which eliminates the requirement to specifically estimate the mass and center of mass and yet produces consistent sinograms, is presented in [41].

Another potential area of future research concerns the coefficient κ, which is used in the sinogram restoration algorithm to indicate the confidence in the given sinogram segmentation. A larger value indicates a higher degree of confidence, so a very large value of κ could be

used if the true segmentation were known and a smaller value for estimated segmentations. It may be reasonable to let κ vary spatially to account for our varying degrees of confidence in the sinogram segmentation. In particular where there was significant interpolation, we would expect to make κ smaller. Also, one might want κ to be small near the estimated support value and increase with increasing t. The rate of increase might be related to the performance measure of the support value estimation algorithm. A heuristic approach has been reported in [42].

Finally, another area of future research would consider the possibility of spatially varying coefficients β and γ. These coefficients specify the expected spatial smoothness of sinograms, which one might expect to be related to the spatial smoothness and shape of objects. Some results along these lines have been reported [43].

References

[1] G. S. Harell, D. F. Guthaner, R. S. Breiman, C. C. Morehouse, E. J. Seppi, W. H. Marshall, and L. Wexler. Stop-action cardiac computed tomography. *Radiology*, 123:515–517, 1977.

[2] W. H. Rowan et al. Algorithms for limited-angle computed tomography. In *Proc. International Workshop on Physics and Engineering an Medical Imaging*, pages 167–177, 1982. Pacific Grove, CA: IEEE Computer Society.

[3] C. K. Zoltani, K. J. White, and R. P. Kruger. Results of feasibility study on computer assisted tomography for ballistic applications. Technical report, U.S. Army Ballistic Research Laboratory, 1983. ARBRL-TR-02513.

[4] G. T. Herman. *Image Reconstruction from Projections*. Academic Press, New York, 1980.

[5] D. R. Wehner. *High Resolution Radar*. Artech House, Inc., Norwood, MA, 1987.

[6] M. Ein-Gal. *The Shadow Transformation: An Approach to Cross-Sectional Imaging*. PhD thesis, Stanford University, Dept. of Electr. Engr., 1974.

[7] A. K. Louis. Picture reconstruction from projections in restricted range. *Math. Meth. in the Appl. Sci.*, 2:209–220, 1980.

[8] M. E. Davison and F. A. Grunbaum. Tomographic reconstructions with arbitrary directions. *Comm. Pure Appl. Math.*, 34:77–119, 1979.

[9] T. Inoye. Image reconstruction with limited angle projection data. *IEEE Trans. Nucl. Sci.*, NS-26(2):2666–2669, 1979.

[10] J. A. Reeds and L. A. Shepp. Limited angle reconstruction in tomography via squashing. *IEEE Trans. on Medical Imaging*, MI-6(2):89–97, June 1987.

[11] S. L. Wood, A. Macovski, and M. Morf. Reconstruction with limited data using estimation theory. In Raviv et. al., editor, *Computer Aided Tomography and Ultrasonics in Medicine*, pages 219–233. North-Holland Publishing Co., 1979.

[12] M. H. Buonocore. *Fast Minimum Variance Estimators for Limited Angle Computed Tomography Image Reconstruction*. PhD thesis, Stanford University, 1981.

[13] M. I. Sezan and H. Stark. Tomographic image reconstruction from incomplete view data by convex projections and direct Fourier inversion. *IEEE Trans. Med. Imag.*, MI-3(2):91–98, 1984.

[14] B. P. Medoff. Image reconstruction from limited data: theory and applications in computerized tomography. In H. Stark, editor, *Image Recovery: Theory and Application*, chapter 9, pages 321–368. Orlando: Academic Press, 1987.

[15] R. W. Gerchberg. Super-resolution through error energy reduction. *Optical Acta*, 21(9):709–720, 1974.

[16] A. Papoulis. A new algorithm in spectral analysis and band-limited extrapolation. *IEEE Trans. Circuits Syst.*, CAS-22(9):735–742, 1975.

[17] J. L. C. Sanz and T. S. Huang. Unified Hilbert space approach to iterative least-squares linear signal restoration. *J. Opt. Soc. Am.*, 73(11):1455–1465, 1983.

[18] K. M. Hanson and G. W. Wecksung. Bayesian approach to limited-angle reconstruction in computed tomography. *Appl. Optics*, 24:4028–4039, December 1980.

[19] S. Geman and D. E. McClure. Bayesian image analysis: and application to single photon emission tomography. Technical report, Brown University, 1985. Preprint to appear in 1985 Proc. Amer. Stat. Assoc. Statistical Computing.

[20] D. J. Rossi and A. S. Willsky. Reconstruction from projections based on detection and estimation of objects–parts I and II: Performance analysis and robustness analysis. *IEEE Trans. ASSP*, ASSP-32(4):886–906, 1984.

[21] M. Soumekh. Binary image reconstruction from four projections. In *Proceedings of the 1988 Int'l Conf. Acoust. Speech. Sig. Proc.*, volume 2, pages 1280–1283, April 1988.

[22] S. K. Chang and G. L. Shelton. Two algorithms for multiple-view binary pattern reconstruction. *IEEE Trans. Sys. Man and Cyber.*, pages 90–94, January 1971.

[23] P. C. Fishburn, J. C. Lagarias, J. A. Reeds, and L. A. Shepp. Sets uniquely determined by projections on axes I. continuous case. *SIAM Journal on Applied Mathematics*, 50(1):288–306, February 1990.

[24] H. J. Trussell, H. Orun-Ozturk, and M. R. Civanlar. Errors in reprojection methods in computerized tomography. *IEEE Trans. Medical Imaging*, MI-6(3):220–227, 1987.

[25] J. H. Park, K. Y. Kwak, and S. B. Park. Iteractive reconstruction-reprojection in projection space. *Proc. IEEE*, 73(6):1140–1141, June 1985.

[26] J. H. Kim, K. Y. Kwak, S. B. Park, and Z. H. Cho. Projection space iteration reconstruction-reprojection. *IEEE Trans. Med. Imag.*, MI-4(3), September 1985.

[27] M. I. Sezan and H. Stark. Image restoration by convex projections in the presence of noise. *Applied Optics*, 22(18):2781–2789, 1983.

[28] D. Ludwig. The Radon transform on Euclidean space. *Comm. Pure Appl. Math.*, 19:49–81, 1966.

[29] S. Helgason. *The Radon Transform*. Birkhauser, Boston, MA, 1980.

[30] J. L. Prince and A. S. Willsky. A geometric projection-space reconstruction algorithm. *Linear Algebra and Its Applications*, 130:151–191, 1990. Special Issue on Linear Algebra in Computed Tomography, ed. G. T. Herman.

[31] J. L. Prince and A. S. Willsky. Reconstructing convex sets from support line measurements. *IEEE Trans. Patt. Anal. Mach. Intell.*, 12(4):377–389, April 1990.

[32] J. L. Prince and A. S. Willsky. Convex set reconstruction using prior geometric information. *CVGIP: Graphical Models and Image Processing*, 53(5):413–427, 1991.

[33] S. R. Deans. *The Radon Transform and Some of Its Applications.* John Wiley and Sons, New York, 1983.

[34] M. Ravichandran and F. C. Gouldin. Reconstruction of smooth distributions from a limited number of projections. *Applied Optics*, 27(19):4084–4097, October 1988.

[35] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging.* IEEE Press, New York, 1988.

[36] J. L. Prince. *Geometric Model-Based Estimation From Projections.* PhD thesis, Massachusetts Institute of Technology, 1988.

[37] D. P. Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods.* Academic Press, New York, 1982.

[38] G. Wahba. *Spline Models for Observational Data.* SIAM, Philadelphia, 1990.

[39] A. S. Willsky and H. L. Jones. A generalized likelihood ratio approach to state estimation in linear systems subject to abrupt changes. Technical Report LIDS-P-538, M.I.T. Laboratory for Information and Decision Systems, 1974.

[40] A. M. Mier-Muth and A. S. Willsky. A sequential method for spline approximation with variable knots. Technical Report ESL-P-759, M.I.T. Electronic Systems Laboratory, 1977.

[41] J. L. Prince and A. S. Willsky. Constrained sinogram restoration for limited-angle tomography. *Optical Engineering*, May 1990.

[42] J. L. Prince. Consistency and convexity in object reconstruction from projections. In *Proc. of the IEEE Int'l Conf. Systems Engr.*, 1990. Pittsburgh, PA, August.

[43] J.L. Prince. An iterative approach to sinogram restoration. In *Proc. of the 1990 IEEE Int'l Conf. on Engr. in Med. and Biol.* IEEE Press, 1990.

Figure Captions

Figure 1. The geometry of the 2-D Radon transform.

Figure 2. (a) The M I T ellipse and (b) its sinogram.

Figure 3. The convex support of an object and the support of a projection.

Figure 4. A region of support for the 2-D Radon transform.

Figure 5. Block diagram of hierarchical reconstruction algorithm.

Figure 6. A projection modeled as a linear spline with knots.

Figure 7. Generalized likelihood ratio threshold selection curve.

Figure 8 Support value estimates using the knot-location algorithm shown using diamond markers along with the true support values shown using vertical dotted lines: head-on noisy projections of an ellipse.

Figure 9 Support value estimates using the knot-location algorithm shown using diamond markers along with the true support values shown using vertical dotted lines: broadside noisy projections of an ellipse.

Figure 10 Knot-location support value estimation followed by maximum likelihood support vector estimation for (a) 100.0dB, (b) 10.0dB, (c) 3.0dB, and (d) 0.0dB sinograms.

Figure 11. A 3.0dB noisy sinogram and its reconstruction using convolution backprojection.

Figure 12. Full hierarchical sinogram estimates and object reconstructions for two segmentations.

Figure 13. Objects reconstructed using convolution backprojection applied to (a) the left-most 40 projections, (b) the right-most 40 projections, (c) 15 sparse projections, and (d) 10 sparse projections of a 10.0dB noisy sinogram.

Figure 14. Sinograms restored using the hierarchical algorithm applied to (a) the left-most 40 projections, (b) the right-most 40 projections, (c) 15 sparse projections, and (d) 10 sparse projections.

Figure 15. Objects reconstructed from restored sinograms in corresponding panels of Figure 15.

Figure 16. Two-disk object and noise-free sinogram.

Figure 17. Two limited-angle cases showing CBP reconstruction (left) and hierarchical reconstruction (right).

Figure 18. Full hierarchical results for limited-angle studies using ellipse-based support vector estimation.

Figure 1. The geometry of the 2-D Radon transform
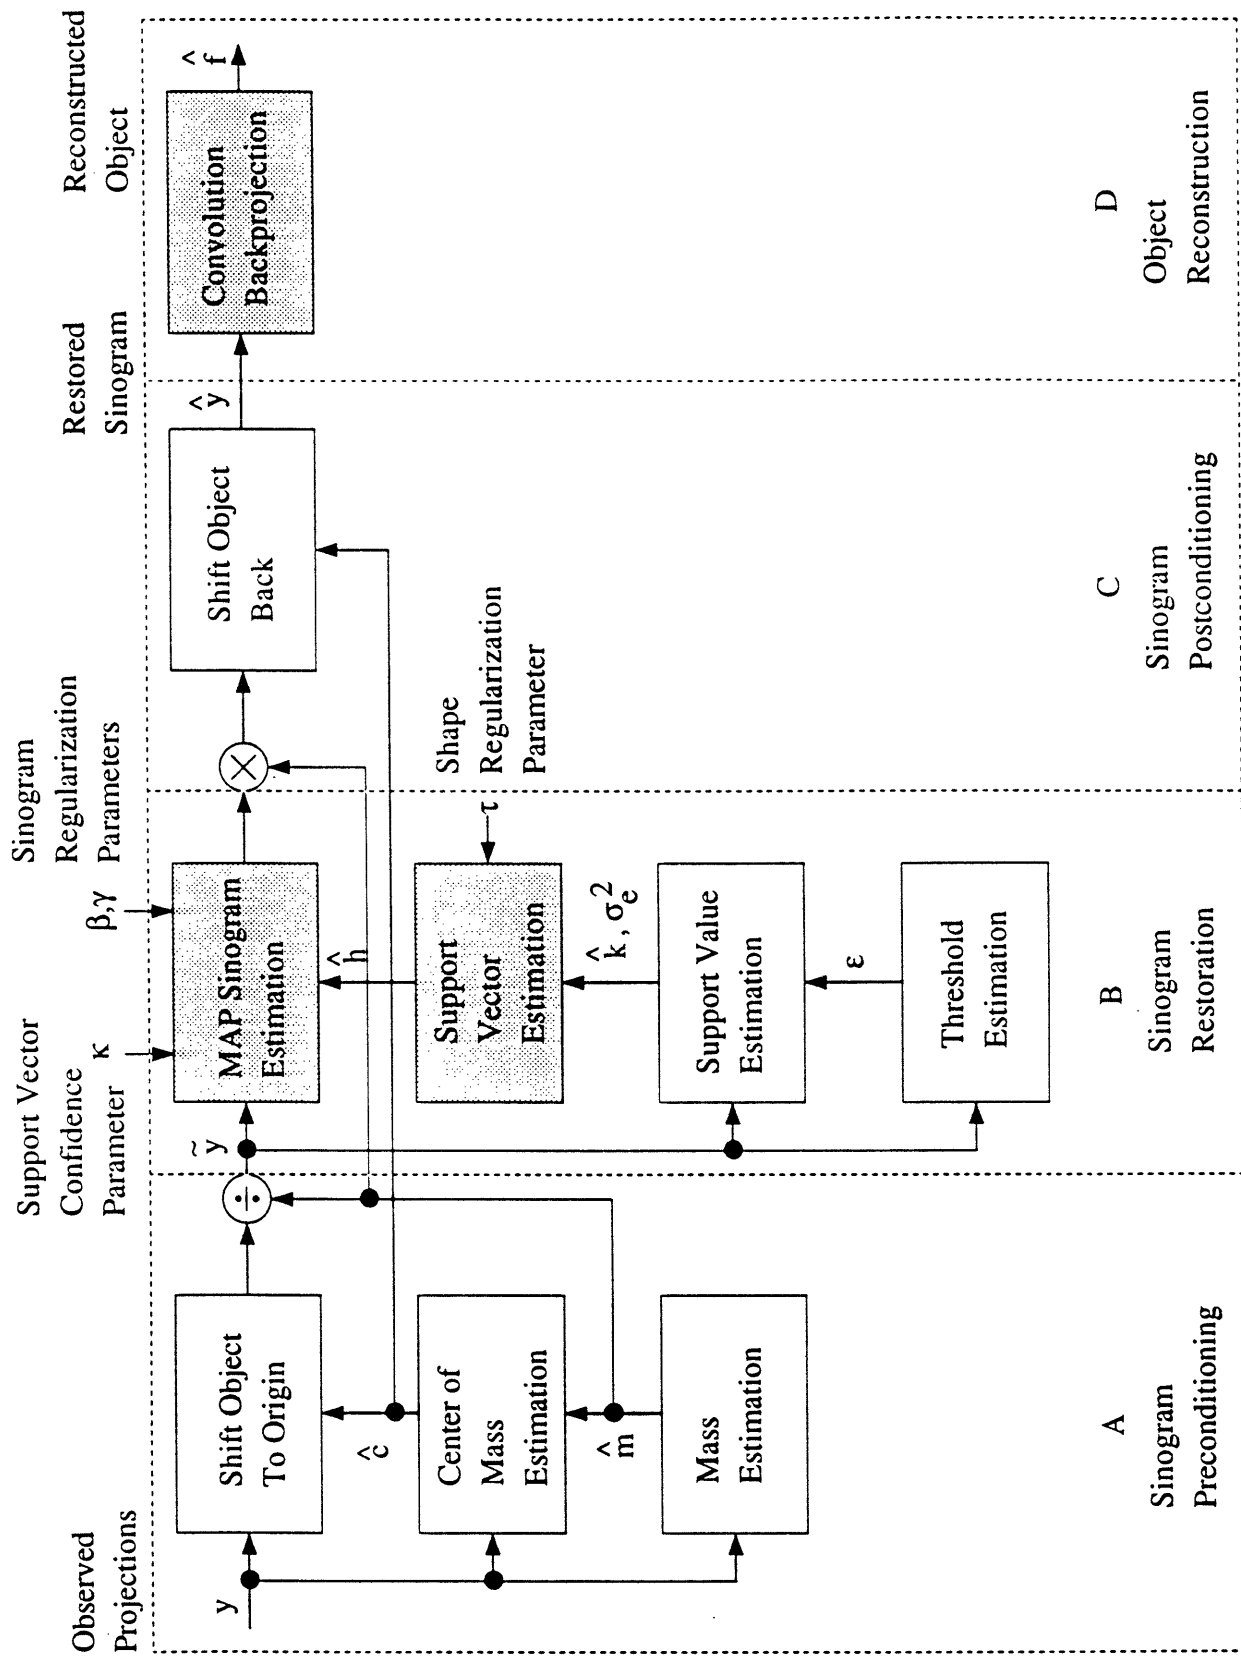
Figure 2(a). The M I T ellipse

Figure 2(b). The sinogram of the M I T ellipse

Figure 3. The convex support of an object and the support of a projection

**Figure 4.** A region of support for the 2-D Radon transform

Figure 5. Block diagram of hierarchical reconstruction algorithm

Figure 6. A projection modeled as a linear spline with knots

Figure 7. Generalized likelihood ratio threshold selection curve.

Figure 8 Support value estimates using the knot-location algorithm shown using diamond markers along with the true support values shown using vertical dotted lines: head-on noisy projections of an ellipse

Figure 9. Support value estimates using the knot-location algorithm shown using diamond markers along with the true support values shown using vertical dotted lines: broadside noisy projections of an ellipse

Figure 10. Knot-location support value estimation followed by maximum
likelihood support vector estimation for (a) 100.0dB, (b) 10.0dB
(c) 3.0dB, and (d) 0.0dB sinograms

(a)   (b)

Figure 11. A 3.0dB noisy sinogram and its reconstruction using convolution backprojection

Figure 12. Full hierarchical sinogram estimates and object reconstruction for two segmentations

Figure 13. Objects reconstructed using convolution backprojection applied to (a) the left-most 40 projections, (b) the right-most 40 projections, (c) 15 sparse projections, and (d) 10 sparse projections

Figure 14. Sinograms restored using the hierarchical algorithm applied to (a) the left-most 40 projections, (b) the right-most 40 projections, (c) 15 sparse projections, and (d) 10 sparse projections

Figure 15. Objects reconstructed from restored sinograms in corresponding panels of Figure 15

(a)

(b)

Figure 16. Two-disk object and noise-free sinogram

Figure 17. Two limited-angle cases showing CBP reconstruction (left) and hierarchical reconstruction (right)

Figure 18. Full hierachical results for limited-angle studies using ellipse-based support vector estimation